



Editorial

Knowledge engineering and ontologies for autonomous systems 2004 AAAI Spring Symposium

1. Problem statement

Many researchers feel that an autonomous system must have an internal representation of entities, events, and situations that it perceives in the world in order for it to behave appropriately in uncertain environments. It must have an internal model that captures the richness of what it knows and learns, and a mechanism for computing values and priorities that enables it to choose effective actions [1]. The term “autonomous system” in this context refers to an embodied intelligent system that can operate for extended periods of time without human supervision. A major challenge for these systems is maintaining an accurate internal representation of pertinent information about the environment. The inability to do this well hinders effective task planning and execution.

Although a large body of work exists in various knowledge representation, ontology, and data fusion areas, relatively little has been applied to real-time world modeling in autonomous systems. Research in autonomous systems has reached a level of maturity such that it could greatly benefit from leveraging the work that has been on-going in these areas. World modeling in autonomous systems provides a rich context in which to apply theoretical and practical knowledge representation, ontological, and data fusion techniques.

The field of autonomous systems continues to gain traction both with researchers and practitioners. Funding for research in this area has grown over the past few years, and recent high profile funding opportunities have started to push theoretical research efforts into

practical use. However, much research still needs to be performed in the area of knowledge representation, a *vital component of many autonomous systems*.

2. The symposium

The Knowledge Representation and Ontologies for Autonomous Systems Symposium was motivated by the desire to bring together experts in the autonomous systems, knowledge representation, ontology, and data fusion communities to explore leveraging existing knowledge technologies to benefit autonomous systems. It was held during March 22–24, 2004 at the Stanford Campus in Palo Alto, CA as part of the 2004 American Association for Artificial Intelligence (AAAI) Spring Symposium Series. The symposium was the first on this topic, and was attended by participants representing a cross-section of the communities mentioned above. The primary goals of this symposium were threefold: (1) to educate the autonomous systems community as to the strengths and weaknesses of various knowledge representation approaches, (2) to educate the knowledge representation community as to the knowledge-related challenges being faced within the autonomous systems arena and (3) establish networks of teaming arrangements and possible collaborations to allow the communities to work closer together in the future. Detailed objectives included:

- (1) exploring how knowledge representation technologies can be used to capture and reason

- with parametric, spatial, dynamic and symbolic knowledge;
- (2) exploring the usefulness of different types of ontologies for autonomous systems;
 - (3) exploring the best ways of representing a priori and in situ knowledge, value judgments, state information, history, plans, entities, events, situations, intent, task knowledge, and self-knowledge;
 - (4) determining which knowledge technologies work best for different challenges in autonomous systems, including corresponding performance measures;
 - (5) understanding the requirements that subsystems (e.g., sensors, learning modules, planners, and operator control units) place on knowledge representations;
 - (6) understanding and formalizing the interaction between disparate knowledge representations (e.g., images, maps, classes, and relationships) that provide complementary information about the same object or event;
 - (7) understanding the role of knowledge in model-based perception and control;
 - (8) identifying approaches to formalizing the autonomous system's internal representation;
 - (9) means to measure the quality of knowledge within autonomous systems;
 - (10) exploring the reusability of knowledge among disparate autonomous systems;
 - (11) determining how data fusion technologies (which support autonomous system sensing capabilities) can be assisted by using knowledge technologies;
 - (12) identifying mechanisms to ensure a tightly collaboration between colleagues in the autonomous systems and knowledge technology communities.

The symposium had two main components: presentations, and a challenge problem. The symposium started with an autonomous systems keynote presentation by Prof. Ernst Dickmanns (University of the Federal Armed Forces of Germany, Munich) who described the state of the art in autonomous vehicle research and development, focusing on efforts over the past 25 years at his university. They have been successfully driving autonomous vehicles on public roads for 12 years; a key to their success is the use of specialists for recognizing particular kinds of objects expected

in an on-road environment. This was followed by six paper presentations that were grouped into the tracks "Knowledge Representation Perspectives and Integration Issues" and "Knowledge Representation for Autonomous Mobility". A poster session followed the paper presentations.

On the second day, Dr. Michael Genesereth of Stanford University gave a knowledge representation keynote presentation entitled "World Models for Autonomous Systems" in which he described his thoughts on what types of knowledge representations appeared to provide the most value to autonomous systems. He argued that state machine representations are limited, there is a need for more extensible approaches using logic and probability. In the longer term, robots need the ability to dynamically reformulate its world model, guided by its purposes. This was followed by four paper presentations in the track "Applying Ontologies to Autonomous Systems".

Some questions seemed to be common following many of the presentations. They were: "How does one know what knowledge should be embedded in an external knowledge base versus in the code itself?", "How does one know which representations are good for what types of requirements?", and "What is an ontology and how is it different from the knowledge representation techniques we have used in the past?". The answers to the first two questions varied from presenter to presenter, showing that there is no clear-cut answer and more research needs to be performed. To address the third question, Michael Uschold from Boeing started off the second day with a brief presentation giving an overview of what ontologies are and some of their primary uses.

After the presentations, the audience was split into three predefined, crossdisciplinary breakout groups, each tasked with addressing the challenge problem. Their job was to determine a "knowledge architecture" for a group of five trash-removal robots that were responsible for cleaning an airport. The robots had to coordinate with each other, provide complete trash-removal coverage of the airport multiple times each day, monitor their own health, travel within marked lanes whenever possible, recycle, identify suspicious packages, and stay a predefined distance from humans at all times. Within the "knowledge architecture", the groups had to define the types of knowledge necessary for the robots to perform their tasks, identify the types of representations that lent themselves best to

representing that type of information, and develop the interfaces between the knowledge sources and the algorithms that were controlling the robots.

On the third day, the moderators of the breakout groups reported back on their groups' findings. As expected, the three groups chose different approaches to tackle the problem with two of the groups focusing on creating a task decomposition, and the third focusing on creating a knowledge decomposition. This provides insight that there is no "magic bullet" in knowledge representations, and that different techniques offer different advantages and disadvantages. Allowing the participants to "get their hands dirty" by addressing the challenge problem also confirmed the belief that knowledge representation for autonomous systems is a tough problem, and should receive more attention from the community. This challenge problem and the breakout groups approaches to address it is described in further detail in a paper entitled "Knowledge Representation For A Trash Collecting Robot: Results From The 2004 AAAI Spring Symposium" within this journal issue.

We concluded the workshop with a panel discussion made up of participants from industry and academia, and representing all of the communities present at the symposium. The panel was tasked with highlighting the main issues and challenges that came out of the symposium, as well as determining the best way for these communities to work together in the future. Issues that arose from the panel included the need for an upper ontology, the challenge of integrating disparate terminology and semantics from different disciplines, and the need for a knowledge representation formalism to capture the autonomous systems' competencies. There was also widespread agreement that the symposium was valuable and that similar ones should be held in the future.

3. Results and future direction

This symposium was intended to be the first in a series of workshops that address the general area of applying knowledge representation techniques towards the area of autonomous systems. As mentioned earlier in the article, our primary objectives were to bring the two communities together to explore the potential benefits of leveraging knowledge representation technologies to meet challenges in autonomous systems and set

up collaborations. To a large extent, we succeeded in these goals. There is room for improvement for KR researchers to point out the weaknesses as well as the strengths of their approaches. This would make it much easier to make a qualitative comparison of the various approaches. Future workshops participants will be encouraged to give balanced assessments of the pros and cons of their approaches.

As was evident throughout the course of the symposium, there are many different approaches to representing knowledge in autonomous systems, with no clear-cut "winner". This initial symposium provided an opportunity to gain a better understanding of the capabilities of various knowledge representation techniques as well as a deeper understanding of the challenges being faced by the autonomous systems community. It also made participants aware of the large number of approaches that are available for representing knowledge within autonomous systems. By allowing these two communities to come together focusing on a common cause, it is the hope that a stronger joint community will be formed, that will be able to apply the theoretical capabilities of knowledge representation approaches with the practical challenges being faced in the world of autonomous systems.

Towards this end, future workshops will focus the group on more specific challenges. A number of functional areas were highlighted throughout the symposium that showed a profound need for formal approaches to knowledge representation. They include robot localization, predictive models for moving object prediction, cost-based and state machine-based planning approaches, and system integration. Future workshops will focus on a subset of these areas, with an emphasis on determining which knowledge representation techniques are more appropriate for different classes of knowledge. Challenge problems will continue to be used to keep the group grounded in real-world problems, which the scope of the challenge problem much more specific and focused in the functional area of interest.

4. Synopsis of this special issue

This special issue has been organized to ensure that the significant results presented at the symposium reach a wider audience. We have asked selected authors from

the symposium to submit updated and extended versions of their paper for inclusion in this issue of the journal.

The nine articles in this journal issue represent a good sampling of the types of knowledge representation approaches that are currently being applied to autonomous systems. The first set of three articles explores different perspective on knowledge representation with an emphasis on integration issues. Cassimatis, Trafton, Bugajska, and Schultz propose an architecture that integrates disparate reasoning, planning, sensation, and mobility algorithms by composing them from strategies for managing mental simulations. Wagner, Visser, and Herzog propose an egocentric representation which relies on 1D ordering information that still provides sufficient allocentric information to solve navigation and localization tasks. Potts and Hengst present an algorithm that avoids any protracted period of initial exploration by discovering multiple levels of a task hierarchy simultaneously.

The next set of five articles explores applying knowledge representation techniques towards autonomous mobility issues. Balakirsky and Scrapper provide an overview of a real-time deliberative planning system and describe the areas that knowledge has been applied to limit the system's graph complexity. Barbera, Albus, Messina, Schlenoff, and Horst explore the 4D/RCS methodology, and how detailed task knowledge is represented in a task context-sensitive relationship that supports very complex real-time processing that a computer control system must perform. Wood presents the view that the reliability of represented knowledge guides information seeking and perhaps explains why some things get ignored. Henninger and Madhavan compares the performance of an extended Kalman filter based model, a neural net based model and a Newton based dead-reckoning model, all used to predict an agent's trajectory and position. Reichard describes a generic architecture for the implementation of health monitoring within a complex system, the representation of system health information, and an approach for integrating health information with autonomous control.

The final set of two papers explore the application of ontologies to autonomous systems. Wray, Lisse, and Beard argue that combining ontology representation and agents optimized for performance can capitalize on the strengths of individual approaches and reduce

individual weaknesses. Provine, Schlenoff, Balakirsky, Smith and Uschold report on the results of a first implementation demonstrating the use of an ontology to support reasoning about obstacles to improve the capabilities and performance of on-board route planning for autonomous vehicles.

Acknowledgements

We would like to thank the symposium's organizing committee who helped to make the event a success by providing useful technical insights on the topics for the symposium and in assisting in reviewing the papers that were submitted. The organizing committee consisted of:

- (1) Craig Schlenoff (chair), National Institute of Standards and Technology (NIST), USA
- (2) Michael Uschold (co-chair), Boeing, USA
- (3) James Albus, NIST, USA
- (4) Stephen Balakirsky, NIST, USA
- (5) James Crawford, NASA Ames Research Center, USA
- (6) Hugh Durrant-Whyte, The University of Sydney, Australia
- (7) Michael Gruninger, University of Maryland, College Park, USA
- (8) Otthein Herzog, University of Bremen, Germany
- (9) Benjamin Kuipers, University of Texas at Austin, USA
- (10) Elena Messina, NIST, USA
- (11) Illah Nourbakhsh, Carnegie Mellon University, USA
- (12) Charles Shoemaker, US Army Research Laboratory, USA

We would also like to thank the authors who have contributed their updated papers. Finally, we would like to thank Professor Thomas Henderson and the team at Elsevier for the kind advice and help to complete this special issue on the Knowledge Representation and Ontology for Autonomous Systems Symposium.

Reference

- [1] J. Albus, A. Meystel, *Engineering of Mind*, Wiley, 2001.



Craig Schlenoff received his Bachelors degree in mechanical engineering from the University of Maryland, College Park and his Masters degree in mechanical engineering from Rensselaer Polytechnic Institute. He is a researcher in the Intelligent Systems Division at the National Institute of Standards and Technology. His research interests include knowledge representation, ontologies, and process specification, primarily applied to autonomous systems and manufacturing.

He has recently served as the program manager for the Process Engineering Program at NIST as well as the Director of Ontologies and Domain Knowledge at VerticalNet, Inc. He has served on the organizing and program committee of numerous knowledge representation and ontology related conferences and workshops, and has published over 35 papers in related areas.

Craig Schlenoff
*Intelligent Systems Division, National Institute of
 Standards and Technology, 100 Bureau Drive
 Stop 8230, Gaithersburg, MD 20899, USA*
E-mail address: craig.schlenoff@nist.gov



Mike Uschold is a research scientist at Boeing Phantom Works, the advanced research and development organization of The Boeing Company. His interests center around the field concerned with the development and application of ontologies. This includes the emerging Semantic Web, semantic integration, knowledge management, and more recently, in the area of world modeling for autonomous vehicle navigation. For over two decades, Mike has been involved in a

wide range of activities in these areas, including research, applications and teaching. Dr Uschold is on the industrial advisory boards of various projects and initiatives related to the Semantic Web and other knowledge technologies. He is very active in organizing and participating in workshops and conferences on these topics. He received his BS in mathematics and physics at Canisius College in Buffalo, N.Y. in 1977, a Masters in computer science from Rutgers University in 1982, and a PhD in Artificial Intelligence from The University of Edinburgh in 1991. Before arriving at the Boeing Company in 1997, Dr Uschold was a senior member of technical staff in the Artificial Intelligence Applications Institute (AIAI) at the University of Edinburgh. He has also been a lecturer and a research associate at the Department of AI at the University of Edinburgh.

Michael Uschold
*Boeing, Phantom Works, P.O. Box 3707
 mls 7L-40, Seattle, WA 98124, USA*
E-mail address: michael.f.uschold@boeing.com